




US 20020188735A1

(19) **United States**(12) **Patent Application Publication** (10) **Pub. No.: US 2002/0188735 A1****Needham et al.**(43) **Pub. Date: Dec. 12, 2002**(54) **PARTIALLY REPLICATED, LOCALLY
SEARCHED PEER TO PEER FILE SHARING
SYSTEM****Publication Classification**(51) **Int. Cl.⁷ G06F 15/16**(52) **U.S. Cl. 709/229**(76) **Inventors: Bradford H. Needham, North Plains,
OR (US); Mark Lewis, La Grande, OR
(US)****Correspondence Address:****BLAKELY SOKOLOFF TAYLOR & ZAFMAN
12400 WILSHIRE BOULEVARD, SEVENTH
FLOOR
LOS ANGELES, CA 90025 (US)**(57) **ABSTRACT**

Controlling access to a file includes selecting, by a first network node in a peer to peer file sharing system, a file to make available for sharing with a second network node in the peer to peer file sharing system, the file being stored on the first network node. The first network node sends index information associated with the selected file to the second network node. The second network node subsequently performs a local search of a set of index information associated with a set of files stored in the peer to peer file sharing system to locate the index information for the selected file; and requests a copy of the selected file from the first network node using the index information.

(21) **Appl. No.: 09/878,023**(22) **Filed: Jun. 6, 2001**




Figure 1

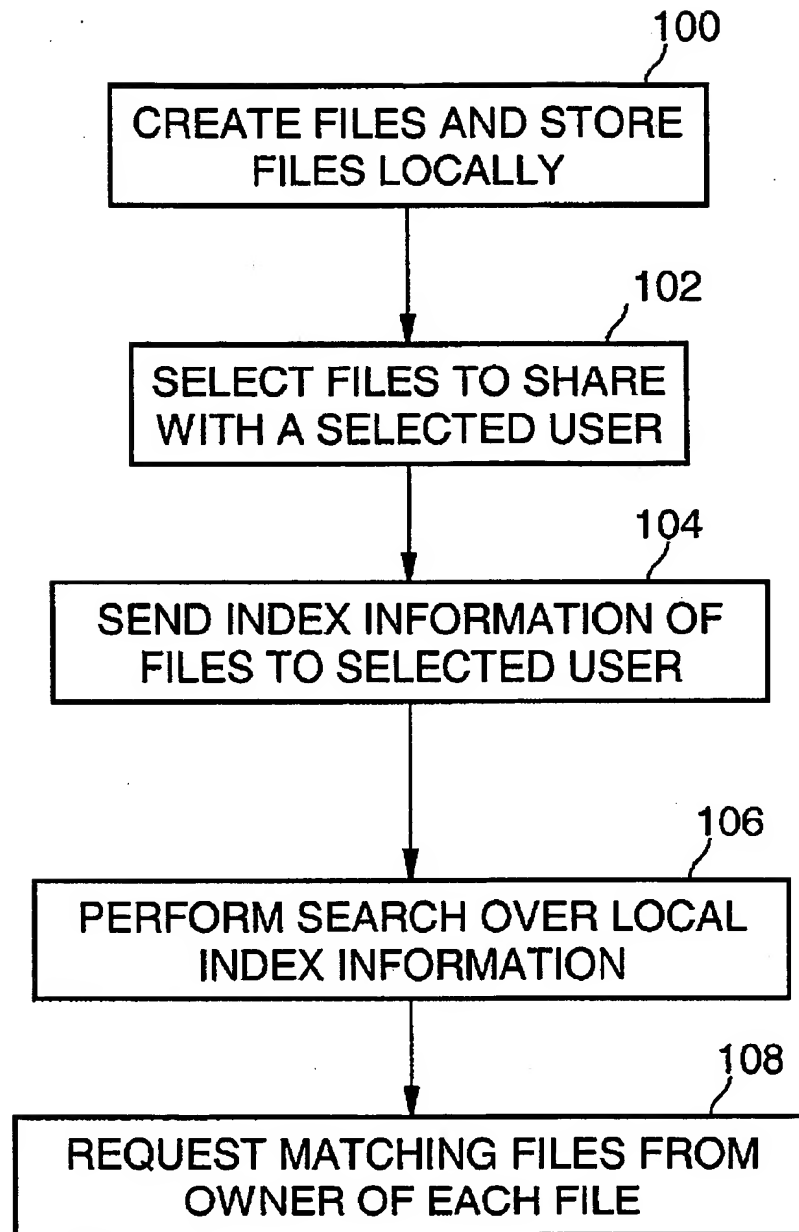


Figure 2

PARTIALLY REPLICATED, LOCALLY SEARCHED PEER TO PEER FILE SHARING SYSTEM

BACKGROUND

[0001] 1. FIELD

[0002] The present invention relates generally to distributed file systems and protocols and, more specifically, to file security in peer to peer media sharing systems.

[0003] 2. DESCRIPTION

[0004] There are at least several peer to peer media sharing systems in current use on computer networks such as the Internet. Napster™ is a popular system wherein a central server performs all searches for requested files and returns the names of files and their network addresses on each system in the network that match the request. In Gnutella™, another popular peer to peer system, users make search requests that are propagated throughout a network and perform searches on each machine in the network. In FreeNet™, there is no search ability. To retrieve a file using FreeNet, the user must already know the file's true name, which may be extremely difficult to guess.

[0005] Both Gnutella and Napster store index information to support searches outside the requester's machine (e.g., on other machines in the network), causing scarce network bandwidth to be consumed every time any user in the network performs a search for a desired file. Neither Gnutella nor Napster provide any hiding of files in the network. In either system, making a file available to one user of the network makes it available to all users of the network. This may be undesirable for personally-created files, because the owner may wish to make only certain files available to certain people, but not to the world at large. Napster or Gnutella could be modified to associate file access permissions with index information, but this produces other problems. In Gnutella, because index information is replicated as files are shared among network nodes, the system would then have to solve the problem of maintaining consistent permissions in replicated indices. In Napster, after adding permissions to each file in the main database, the system would still require a central server with the associated disadvantages of a lack of fault tolerance and difficulty in maintaining scalability.

[0006] For peer to peer file sharing systems supporting the sharing of personally created media such as digital photographs, videos, and music, the lack of privacy inherent in these systems is troubling to some users. Furthermore, the search capabilities of such systems are inefficient.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The features and advantages of the present invention will become apparent from the following detailed description of the present invention in which:

[0008] FIG. 1 is a diagram of an example portion of a network illustrating controlled file sharing according to an embodiment of the present invention; and

[0009] FIG. 2 is a flow diagram of file sharing processing according to an embodiment of the present invention.

DETAILED DESCRIPTION

[0010] An embodiment of the present invention is a method of controlling sharing and searching of file infor-

mation in a distributed file system operating as a private peer to peer network. In one embodiment, the invention may be implemented in a peer to peer media sharing system. In typical peer to peer media sharing systems such as Napster and Gnutella, the file data being shared among users is professionally created and widely distributed data (e.g., music files encoded in the Moving Pictures Experts Group (MPEG) audio layer 3 (MP3) format, digital video clips encoded in various formats, etc.). That is, the data may be posted and made available in multiple copies and at multiple network nodes within the distributed network. In contrast, in embodiments of the present invention, the system may be designed to selectively share individually created and narrowly distributed data. For example, the data being shared may be digital photographs or digital videos created by an individual for sharing only among his or her immediate social connections such as family and friends, but not with others. The present invention takes advantage of the single point of creation of this type of data (and single point of posting the data) to control sharing of the knowledge of the existence of files through a private peer to peer network. By controlling the sharing of the knowledge of the files on the network, privacy may be maintained.

[0011] In embodiments of the present invention, index information about files may be communicated among nodes of the network separately from the file data. The index information may be selectively propagated through a public network such as the Internet only between nodes operated by people having social connections (e.g., friends and family), and the file data may be selectively propagated more narrowly in the network, again only between friends and family. This in effect creates a private peer to peer network within a larger public network (such as the Internet, for example).

[0012] In the present invention, individually created and narrowly distributed media may be shared with selected people such that the media is largely undiscoverable by others. There are many prior art systems where all index information is shared with other parties (such as distributed text retrieval systems), and in which the subsets of index information are chosen for reasons other than privacy (for example, mobile computer synchronization systems). In contrast, embodiments of the present invention share selected subsets of index information in a distributed file system with selected other users in a controlled manner, thereby improving the security of shared files of the peer to peer network and the privacy interests of the users of the ad hoc private peer to peer network. Since in the present invention index information is shared only about files that one user wishes to make available to another user, users who do not receive the index information about a file cannot easily discover that file's existence, and thus cannot request that file from the first user's system. Embodiments of the present invention also perform file searches locally, reducing consumption of network bandwidth, yet preserving advantageous fault tolerance of a configuration not requiring a central server.

[0013] Reference in the specification to "one embodiment" or "an embodiment" of the present invention means that a particular feature, structure or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, the appearances of the phrase "in one embodiment" appearing in

various places throughout the specification are not necessarily all referring to the same embodiment.

[0014] Embodiments of the present invention comprise a network of file storage systems, each file storage system having index information describing one or more of the files on that one system and a (possibly empty) subset of files on other systems in the network. Each file may have a globally unique identifier associated with the file that may be very difficult to guess. The identifier may be part of the index information for the file. Each system has the ability to communicate with a subset of nodes in the network (for example, a user of a file storage system may communicate with his or her friends and family via their file storage systems). In various embodiments, a file storage system may be a personal computer (PC), an engineering workstation, a set top box, a personal digital assistant (PDA), a cellular phone, a mainframe computer, an Internet appliance, or any other device for storing and accessing file data via a communications network. A file storage system may also be referred to herein as a network node. Each network node may communicate with one or more other network nodes. The network may be a local area network (LAN), a wide area network (WAN), the Internet, or any other communications network. The files may store any digitally represented data, for example, audio data, video data, image data, or text data. In one embodiment, the file data comprises digital photographs and the file storage systems (nodes) are PCs coupled to one another via the Internet, although the invention is not limited in scope in this respect. In this embodiment, each PC stores the owner's collection of personally created digital photographs and videos, and photographs and videos that friends or family members have shared with the owner.

[0015] FIG. 1 is a diagram of an example portion of a network for controlled file sharing according to an embodiment of the present invention. This example illustrates multiple network nodes interacting with each other in a privately organized peer to peer network based on social connections. In one embodiment, the private peer to peer network may be formed to communicate digital photographs amongst network members, with each digital photograph being stored in a file, although the data shared may be of any type or for any purpose. In this example, John 10 knows Bob 12 and Mary 14. Of course this example is merely for illustrative purposes and does not limit the scope of the invention. The circles shown in FIG. 1 represent the file storage systems (network nodes) of the people discussed in the example, yet there may be many more network nodes in the network that are not shown.

[0016] Assume for the sake of illustration that John has some personally created media that he wishes to share with selected other people he knows. The files may be stored on John's system 10. As part of creating each file storing the data (e.g., digital photograph data), Bob assigns the file a globally unique and very hard to guess identifier (for example, file #13572431234890), using any of several well known methods for generating globally unique and/or very hard to guess identifiers. Identifiers for the files may be chosen to be very hard to guess in order to reduce the ability of another network node to request copies of files for which the node has not received index information about the files. John decides to make part of his media collection, part A 16, available for access by his friend Bob 12, and another part,

part B 18, available for access by his friend Mary 14. John also decides to keep another part, part C 20, private. In other words, John does not want to make any files in part C of his collection available for access by anyone (at least not to Bob or Mary at this point in time).

[0017] In order to selectively make certain files available for access by others in the private peer to peer network, John's system 10 sends index information 22 describing part A of his collection to Bob's system 12, and index information 24 describing part B of his collection to Mary's system 14. Note that the size of the index information communicated to the other network nodes is relatively small compared to the size of the files (photographs, videos, etc.) to be shared. The index information may comprise data such as, but not limited to, a unique and very hard to guess file identifier, a caption or title, an image thumbnail, a creation date, a creator name, file size, format type, and other identifying information. Of course, FIG. 1 describes merely one simplistic example, and with the present invention John's system could send index information for any possible subset of John's files to any one or more other users in the network, in any combination and at any time. The example shown in FIG. 1 implies that the files in part A are exclusive of the files in part B or part C, but the invention is not limited in scope in this respect; the parts of the file collections specified in any given set of index information sent to one or more network nodes may be overlapping.

[0018] At some time in the future, Bob may search his locally stored index information for some particular file or files. For example, suppose Bob is looking for all photographs taken on June 12 at Bob's birthday party, which was attended by John, Mary, and others. The locally stored index information may include index information about Bob's own files (which are stored on Bob's system), as well as index information about the files of others (which are not stored on Bob's system). The index information about the files of others may have been received from those friends and family members as described immediately above. In this example, the local search may return information about some of Bob's photographs, and about some of John's photographs, from part A 22 of John's collection, the index information for which was previously shared with Bob. Bob decides he wants to see a selected one of John's photographs from the party from John's collection that Bob knows about, so Bob's system 12 retrieves the globally unique identifier for the selected file from the associated index information stored locally on Bob's system. Bob's system then sends a file request 26 for the selected file from part A of John's collection (e.g., for file #42258543274590) to John's system. John's system fields the request and returns the requested file to Bob's system.

[0019] Notice that since Mary does not have index information about part A of John's collection, she cannot easily request any of those files (e.g., photographs), and is unaware that those files exist. Mary can only request files in part B of John's collection, because she only has index information, including the file identifiers, for files in part B. Bob cannot request any files from part B of John's collection, because he is unaware that they exist. Both Bob and Mary are unaware of the existence of part C of John's collection, which he has not shared with either person. Since the globally unique file identifiers are assigned such that they are very hard to guess, it may be very unlikely that Bob can guess the file identifier

of a file in John's collection for which Bob has not been informed of the associated index information. Thus, the present system of selectively sharing index information may improve file security and user privacy. Note also that searches for index information for selected files are performed locally rather than being broadcast to all network nodes. No global searches for files may be allowed. Instead, file requests may be made only for files for which index information has been shared. Hence, the sending of index information operates to grant access to selected files. Each network node searches the index information stored locally to determine what files the node has been given access to on other nodes. In some embodiments, the propagation of the index information through nodes of the network may be limited in order to curtail widespread dissemination of the files.

[0020] FIG. 2 is a flow diagram of file sharing processing according to an embodiment of the present invention. At block 100, a first network node creates one or more files and stores the files locally at the first network node. The files may contain any data. Each file may be assigned a globally unique and very hard to guess identifier. At block 102, one or more files to be shared with a selected user (such as a friend or relative) may be selected. At block 104, index information for each of the files to be shared is created and sent to the selected user (i.e., the selected user's network node or file storage system). At a minimum, the index information includes the globally unique and very hard to guess identifier for the file. The sharing of index information may be performed for any number of network nodes, in any combination of sets of index information. The sharing of the index information may in some embodiments be only with people that have social connections (e.g., they are friends, family members, etc.). The set of index information that may be shared from one network node to another network node may be a subset of the index information for all files available in the peer to peer file sharing system. Blocks 100, 102, and 104 may be performed by a first network node that creates one or more files to be shared. In other embodiments, the index information may be propagated through the network to other nodes.

[0021] At block 106, a second network node that has received index information from another network node (which may or may not be the node that created the index information (e.g., the first network node)) performs a search over all local index information available to that second network node. The index information stored locally at the second network node in effect makes the second node aware of and grants access to files stored on the first network node and possibly other nodes. At block 108, the second network node requests one or more matching files from the owners of each file based on the file identifiers contained in the locally stored index information. Thus, the index information is used to obtain a copy of a selected file.

[0022] The present invention thus provides better privacy than existing peer to peer methods and systems. In Gnutella, for example, one cannot control who queries your publicly accessible file directories. In contrast, with the present invention, access may be granted only to those informed of the file identifiers via the sharing of index information. Blanket queries of file directories for index information or other file information may not be allowed. Instead, only requests for the actual file data may be supported. In some

embodiments, either the files or the index information, or both, may be encrypted using known methods to provide additional security prior to sharing either the index information or the file with another network node.

[0023] The techniques described herein are not limited to any particular hardware or software configuration; they may find applicability in any computing or processing environment. The techniques may be implemented in hardware, software, or a combination of the two. The techniques may be implemented in programs executing on programmable machines such as mobile or stationary computers, personal digital assistants, and similar devices that each include a processor, a storage medium readable by the processor (including volatile and non-volatile memory and/or storage elements), at least one input device, and one or more output devices. Program code is applied to the data entered using the input device to perform the functions described and to generate output information. The output information may be applied to one or more output devices.

[0024] Each program may be implemented in a high level procedural or object oriented programming language to communicate with a processing system. However, programs may be implemented in assembly or machine language, if desired. In any case, the language may be compiled or interpreted.

[0025] Each such program may be stored on a storage medium or device, e.g., compact read only memory (CD-ROM), digital versatile disk (DVD), hard disk, magnetic disk, or similar medium or device, that is readable by a general or special purpose programmable machine for configuring and operating the machine when the storage medium or device is read by the computer to perform the procedures described herein. The system may also be considered to be implemented as a machine-readable storage medium, configured with a program, where the storage medium so configured causes a machine to operate in a specific manner. Other embodiments are within the scope of the following claims.

[0026] While this invention has been described with reference to illustrative embodiments, this description is not intended to be construed in a limiting sense. Various modifications of the illustrative embodiments, as well as other embodiments of the invention, which are apparent to persons skilled in the art to which the inventions pertain are deemed to lie within the spirit and scope of the invention.

What is claimed is:

1. A method of controlling access to a file comprising:

selecting, by a first network node in a peer to peer file sharing system, the file to make available for sharing with a second network node in the peer to peer file sharing system, the file being stored on the first network node;

sending index information associated with the selected file from the first network node to the second network node;

performing, by the second network node, a local search of a set of index information associated with a set of files stored in the peer to peer file sharing system to locate the index information for the selected file; and

requesting, by the second network node, a copy of the selected file from the first network node using the index information.

2. The method of claim 1, wherein the index information comprises a globally unique identifier of the selected file.

3. The method of claim 1, further comprising creating at least one file and associated index information by the first network node.

4. The method of claim 1, wherein the index information is sent from the first network node to the second network node only when a user operating the first network node has a social connection to a user operating the second network node.

5. The method of claim 1, wherein the set of index information stored on the second network node is a subset of the set of index information for all files of the peer to peer file sharing system.

6. The method of claim 1, wherein at least one of the selected file and the index information is encrypted by the first network node.

7. The method of claim 1, wherein reception of the index information grants access to the file associated with the index information.

8. The method of claim 1, further comprising limiting propagation of the index information among network nodes.

9. An article comprising: a storage medium having a plurality of machine readable instructions, wherein when the instructions are executed by a processor, the instructions provide for controlling access to a file by

selecting, by a first network node in a peer to peer file sharing system, the file to make available for sharing with a second network node in the peer to peer file sharing system, the file being stored on the first network node;

sending index information associated with the selected file from the first network node to the second network node;

performing, by the second network node, a local search of a set of index information associated with a set of files stored in the peer to peer file sharing system to locate the index information for the selected file; and

requesting, by the second network node, a copy of the selected file from the first network node using the index information.

10. The article of claim 9, wherein the index information comprises a globally unique identifier of the selected file.

11. The article of claim 9, further comprising instructions for creating at least one file and associated index information by the first network node.

12. The article of claim 9, further comprising instructions for sending the index information from the first network node to the second network node only when a user operating the first network node has a social connection to a user operating the second network node.

13. The article of claim 9, wherein the set of index information stored on the second network node is a subset of the set of index information for all files of the peer to peer file sharing system.

14. The article of claim 9, further comprising instructions for encrypting at least one of the selected file and the index information by the first network node.

15. The article of claim 9, wherein reception of the index information grants access to the file associated with the index information.

16. The article of claim 9, further comprising instructions for limiting propagation of the index information among network nodes.

17. A peer to peer file sharing system comprising:

a plurality of network nodes, wherein each network node creates files and assigns a globally unique identifier to each created file, selects files to be shared with other network nodes, sends index information associated with each of the selected files to at least one other network node, performs local searches of index information obtained from other network nodes and index information created by the network node, and requests copies of one or more of the selected files from another network node using the index information for the one or more selected files obtained in the local searches.

18. The peer to peer file sharing system of claim 17, wherein the files store at least one of digital photographs, digital videos, digital audio data, image data, or text data.

19. The peer to peer file sharing system of claim 17, wherein the network nodes of the peer to peer system are operated by users having a social connection.

* * * * *